White Paper

The New Five 9s:

Achieving Big Iron Storage Availability with Commodity Storage Hardware

By Terri McClure, Senior Analyst

March 2014

This ESG White Paper was commissioned by Nexenta and is distributed under license from ESG.

Contents

High Availability Is Highly Important

Always-on Economy with Exponential Data Growth

The great data deluge is expanding as a wealth of information finds its way online and data is sliced, diced, and accessed around the clock. Most business is conducted online: Knowledge workers are working remotely now more than ever, outside of traditional 9-5 hours. And the always-on, global economy rewards only the strong and well prepared. We may have the technology to house the unprecedented data stores that result from these dynamics, but we've been lacking solutions to make them highly available in an efficient manner. In this increasingly pressurized competitive environment, the Internet of Things is relentlessly creating a global matrix of connectivity; social media is dominating our national psyche; mobility is the new mantra; and downtime is simply *not* an option.

Even sensor data is getting into the act. It may not be getting the attention and notoriety of big data, but it is a major component of the deluge. For example, according to Slashdot, Virgin Atlantic is preparing for a significant increase in data as it embraces the Internet of Things with a new fleet of highly connected planes, each expected to create over half a terabyte of data per flight.[1] This sensor data quickly climbs into the petabyte scale for a single day's worth of flights. Multiply that by weeks, months, and years, and the scale of sensor data gets massive.

The times clearly call for a storage strategy that enables 24/7 data availability at scale. The good news is that one is available. The bad news is that it's the same old big iron storage strategy: If you really need something to be available 24/7, then you have to spend a lot of money and put it on an expensive disk array, and then you have to buy a whole bunch of them. That's the way organizations have been dealing with the storage conundrum. Got a performance problem? Throw hardware at it. Got an availability problem? Just throw hardware at it. This kind of wanton spending should have run its course many years ago, but instead it has been perpetuated by companies that had to spend the money required by these systems in order to remain competitive. The best news is that, thanks to software-defined storage solutions like NexentaStor, there is now an alternative to big iron costs.

IT's Biggest Storage Challenges

It's no wonder that managing data growth and data protection are IT's two biggest storage challenges. In ESG's *2012 Storage Market Survey*, rapid growth and management of unstructured data was the storage challenge cited by the most organizations surveyed, with 40% citing it as a challenge and 15% citing it as their primary challenge.[2]

We're in the early stages of developing cost-effective, software-defined storage solutions that make data highly available. This trend is being driven even more forcefully by the price pressures of storing and protecting unstructured data over time on traditional storage systems, which is reflected in ESG research showing that data protection (39%), hardware costs (39%), and lack of physical space (25%) were the second, third, and fourth most-cited storage challenges respectively.

Some of the other highly ranked answers, including the need to support growing virtual server environments (25%) and the management, optimization, and automation of data placement (19%), line up with the data management, storage, and availability challenges facing today's IT organizations.

Five 9s: The Arbiter of Availability

Any downtime can have a significant impact on the business—in the financial services industry, it can cost millions of dollars per minute. While downtime's hard costs vary significantly by industry, soft costs like lost employee productivity, partner dissatisfaction, and missed business opportunities mean that the expense of downtime can far exceed typical measurements based on transaction volume. And companies may not even be able to measure just

[1] Source: Slashdot, *Boeing 787s to Create Half a Terabyte of Data per Flight*, March 2013.

[2] Source: ESG Research Report, *2012 Storage Market Survey*, November 2012. All other ESG research charts and references in this white paper are taken from this research report, unless otherwise noted.

how far reaching the effects of downtime are. That's why high availability is the "must-have" storage feature reported by the greatest percentage of ESG research respondent buyers (59%) (see Figure 1).

Figure 1. Must-have Storage Features and Capabilities

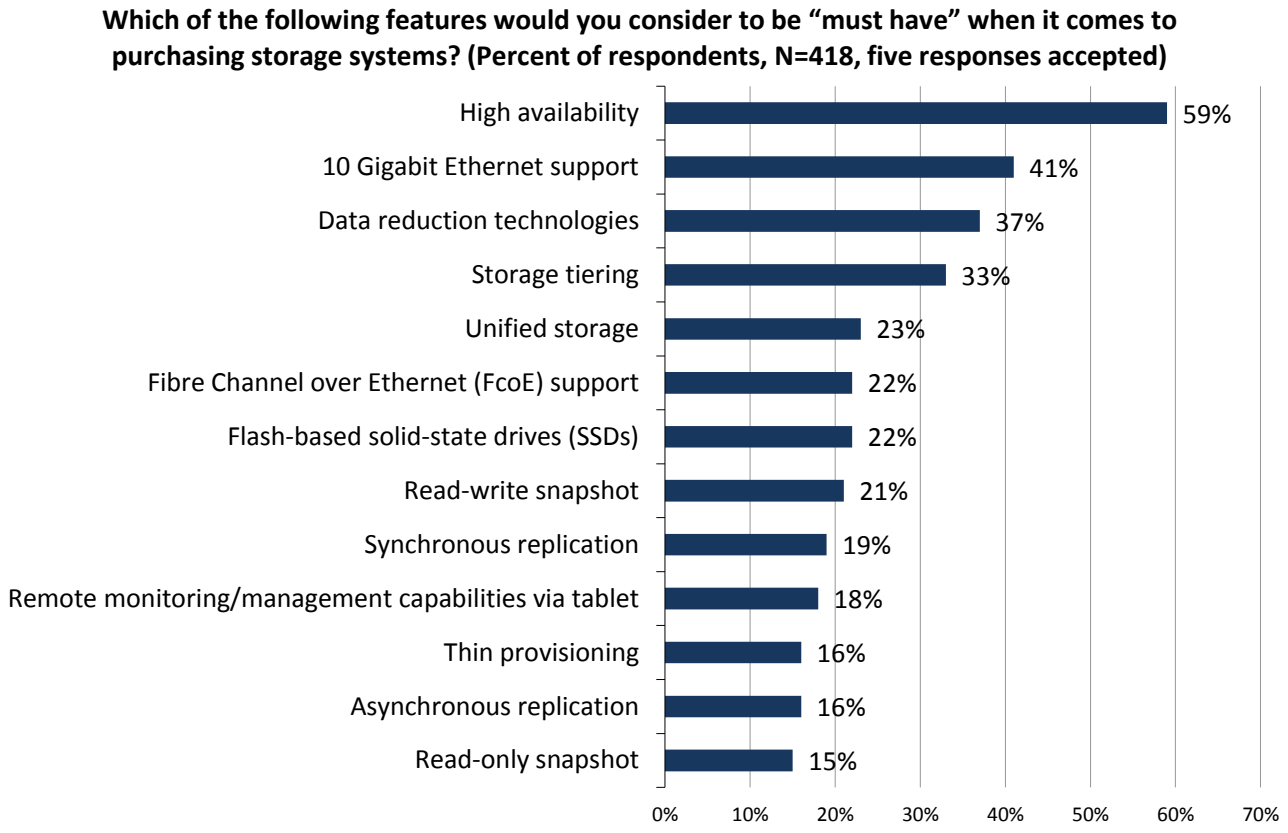

Source: Enterprise Strategy Group, 2014.

The concept of five 9s is a way to measure availability goals as a percentage of system uptime. Five 9s, or 99.999% availability, means the system will be up (or available) 99.999% of the time. The five 9s concept may sound like overkill, but not if you think about what the number means in terms of days and hours of downtime. How much downtime is too much? A system that is 90% available may sound good, *but 90% availability is actually 36.5 days of downtime a year,* and that represents a prohibitive amount of productivity down the drain. Table 1 tells the complete story of 9s availability as it relates to downtime, and when it comes to five 9s—or 99.999% uptime—there is a mere 5.26 minutes of downtime per year, an amount that has become the standard-bearer for availability.

Table 1. Availability Percentage and Corresponding Downtime per Year, Month, and Week

Availability Percentage	Downtime per Year	Downtime per Month	Downtime per Week
90% (one 9)	36.50 days	72 hours	16.8 hours
99% (two 9s)	3.65 days	7.20 hours	1.68 hours
99.9% (three 9s)	8.76 hours	43.80 minutes	10.10 minutes
99.99% (four 9s)	52.56 minutes	4.32 minutes	1.01 minutes
99.999% (five 9s)	5.26 minutes	25.90 seconds	6.05 seconds

Source: Enterprise Strategy Group, 2014.

Every business has different needs, so not all of them require five 9s. Typically, however, mission-critical or business-critical data needs five 9s because it would be disastrous to have an unplanned failure in the middle of a busy day.

Software-defined Storage Changes Everything

Getting to Five 9s with Traditional Approaches

How do you get to five 9s? With five 9s, you need to plan for downtime and protect against a multitude of different failures. In traditional, big iron environments, achieving five 9s requires throwing expensive hardware at the problems so the storage system could withstand and recover from the downtime associated with:

- Hardware failures

  Solution: Disk, controller, disk shelf, and cable recovered with clustering and RAID

- Software failures

  Solution: Snapshot to roll back to a point in time

- Scheduled downtime for maintenance and patches

  Solution: Rolling or online upgrades

- Site failure

  Solution: Remote mirroring to a second, very expensive big iron system

Then there is arguably the most complex source of downtime: the human being. Somebody unplugging a cable by mistake emulates a hardware failure, or somebody inputting an incorrect command emulates a software failure. These breakdowns continue to occur unabated as IT professionals invariably become fatigued, distracted, or otherwise inattentive. It is a safe bet that no matter how well we insulate ourselves from the vagaries of human behavior, human employees will continue to make mistakes. And the storage system needs to have high availability features to deal with them.

The Big Shift from Big Iron

Among its many drawbacks, big iron storage depends on hardware for its intelligence, and the more big iron that companies deploy, the more heterogeneous silos they will end up with at a time when silos represent everything inefficient about the current state of data management (e.g., siloes are typically underutilized). The answer to this dilemma is found by putting a new layer on top of the stack—the software-defined storage (SDS) layer.

Software-defined storage changes the game because it shifts the intelligence to manage availability out of the hardware layer and into a higher order software layer, allowing users to deploy commodity hardware to drive down their CapEx costs by a factor of two to three times and reduce their OpEx as well. Now, the tables are turned, and users can begin riding the commodity curves of the decreasingly expensive hardware market, which allows them to take advantage of lower processor and disk prices.

The situation is even more attractive because in the course of shifting their intelligence from the hardware to the software layer, this essentially virtualizes the back-end of the storage layer with software as well. This makes the system better protected against hardware and other types of failures, ensures that their systems are more highly available, and gives a boost to performance by distributing workloads intelligently.

Benefits of the Software-defined Five 9s Availability Model

ESG believes that we are in a period of rapid change, which is driving higher numbers of new software-defined types of projects at the expense of maintaining legacy systems. Organizations can do almost everything with software-defined storage leveraging commodity hardware that they could do with hardware-based storage—the big difference is that the software-defined approach provides better availability, utilization, and orchestration, while reducing CapEx and OpEx.

Positive Total Cost of Ownership Implications

Total cost of ownership (TCO) has been a major factor in shaping IT budgets and creating cost-effective IT infrastructures for many years. TCO is interpreted differently by CFOs who are focused on managing money, CEOs who have to balance the needs of IT against other enterprise organizations, and CIOs who want to efficiently and effectively utilize all available financial resources in order to provide the strongest possible support for business objectives. All three leaders want to optimize CapEx and OpEx outlays; they are just coming at it from different perspectives.

The relationship between TCO and storage is demonstrated in ESG research, which found that about two-thirds of surveyed organizations identify TCO as one of their five most important criteria for selecting a storage vendor/solution. Figure 2 further validates this concern by showing that 65% of respondents indicated that reducing the cost of storage hardware and software is the most important consideration in their TCO calculation.

Figure 2. Most Important TCO Cost Considerations

Which of the following cost considerations that factor into a TCO calculation would you consider to be the most important? (Percent of respondents, N=271)

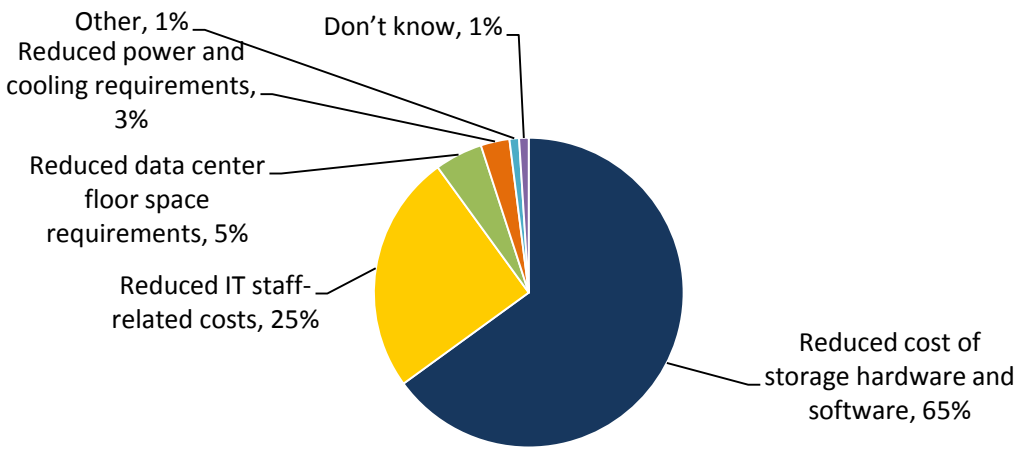

Source: Enterprise Strategy Group, 2014.

Those respondents may be feeling heat from their managers over the size of their next budgets, and from the pressure to do more with the same or fewer financial resources. Software-defined storage solutions allow IT managers to reduce hardware and software costs to accommodate flat or decreasing budgets; thus, five 9s and high availability through SDS may become higher priorities because IT no longer has carte blanche permission to acquire expensive big iron systems whenever they want.

ESG believes that now is a good time for these TCO-conscious IT professionals to embrace software-defined storage by promoting it to their managers not only for its CapEx advantages, but also for its OpEx benefits. Nexenta users report up to a 50% reduction in CapEx thanks to the benefits of running software-defined storage on commodity hardware.

CapEx benefits include:

- **Freedom from specialty hardware.** Users no longer need to pay the tax associated with highly engineered hardware components with embedded specialty controllers and microchips. The intelligence is moved to the software, so users can leverage commodity hardware and hardware upgrades can be performed in a more seamless manner.

- **Vendor-agnostic storage systems.** Users are not tied in to a single hardware vendor and can take advantage of competitively priced commodity hardware.
- **The ability to ride the decreasing commodity pricing curve.** Users drive down storage CapEx over time and stay in line with market price declines as improvements in CPU speeds and disk capacity drive down the cost per GB and IOP.

While users report significant CapEx savings, the OpEx benefits can also be significant with:

- **Enhanced utilization.** In ESG's 2012 storage survey, only 11% of respondents reported that they were completely satisfied with their storage hardware utilization rate. Software-defined storage helps improve hardware utilization because the workload and capacity can be managed across the storage hardware pool.
- **Reduced data center floor space requirements.** Thanks to better capacity management and enhanced utilization, more data is stored on each node, so the overall storage footprint is decreased, saving valuable floor space.
- **Reduced power and cooling requirements.** Another benefit of higher utilization is a reduction in total power and cooling requirements, and their associated cost. Those companies reporting poor hardware utilization rates have an average utilization rate of 52%, which means that almost half of the power and cooling budget is going toward supporting spinning but unused disks.
- **Resource orchestration to ensure workloads get proper hardware allocation and faster provisioning.** This is an important one. Because workload and capacity can be managed at the storage pool level, rather than the individual node level, resources are flexible and can be allocated as needed, significantly speeding time to provision.
- **The ability to keep up with enterprise-wide server and networking virtualization efforts (cloud).** A large majority (95%) of companies that deploy server virtualization run into some sort of storage challenge. The most often reported challenge (43% of server virtualization users surveyed) is the capital cost of the storage infrastructure required to support virtualization. Another 36% report that sizing actual capacity requirements is a challenge, and 28% report that the operational cost of the new storage infrastructure is a challenge. Software-defined storage can help in these areas. Aside from the previously mentioned CapEx and OpEx savings, the orchestration benefits associated with software-defined storage are a better match to virtual environments and can help the storage management team respond to a changing virtual environment in a faster, more efficient manner.

The pitch is simple and succinct: Software-defined storage is the cost-effective wave of the future, and it can be a win-win scenario for all parties.

Nexenta: Enterprise-class Storage Software

Nexenta High Availability—Five 9s with Commodity Hardware

As a major player in the software-defined storage market, Nexenta offers a scalable, node-based architecture managed by software-defined storage at the top layer called NexentaStor. This software-only product features unified block and file storage, provides affordable five 9s availability by protecting from hardware and software failures, enables planned maintenance, and protects users' data. Although Nexenta offers RAID capabilities, it also includes clustering to protect against node failure, along with snapshotting and disk-to-disk backup for data protection from some software failures.

Nexenta's goal is to deliver enterprise-class, hardware-independent storage software that enables organizations to:

- Standardize on a single storage operating system across an ever-increasing set of application workloads via block, file, and object access that is integrated into the user's preferred management framework. Toward that goal, the company offers a single, mature code base, one file system, and one operating system, which are integrated with OpenStack and being developed for additional orchestration frameworks going forward.

- Choose their preferred storage server hardware independently of the storage software, knowing they have the flexibility to change their minds and dual-source or repurpose existing hardware. By offering optimized and certified hardware based on strategic relationships with equipment suppliers, Nexenta also leverages the benefits of an active Illumos community to test, debug, and fix software problems, and provide driver development assistance.
- Deploy a mix of NexentaStor as a bare metal storage solution or virtual storage appliance as business needs dictate. This is supported by the NexentaStor software stack, which can be deployed on a wide range of servers, including Dell, Cisco, HP, SuperMicro, IBM, SGI, Quanta, and others. NexentaStor can also currently be deployed as a virtual storage appliance within VMware, with additional plans for KVM, Xen, and Hyper-V in the works.
- Deploy storage by application with Nexenta VSA for VMware Horizon View, providing accelerated performance, increased desktop density, and auto-configuring and auto-deploying storage—all in a storage vendor-agnostic environment.

Protection from Site Failures

In the case of site failures, Nexenta's MetroHA and Auto-Sync products kick in. MetroHA is a short-hop solution that provides synchronous data mirroring between two sites separated by up to 50 kilometers. Synchronous data mirroring is usually only possible when the nodes utilize a homogeneous software configuration. MetroHA can be configured as either "Active-Active" or "Active-Passive." In Active-Active environments, traffic intended for a failed node is either passed to an existing node or load-balanced across the remaining nodes. Active-Passive provides a fully redundant instance of each node, which is only brought online when its associated node fails. This configuration typically requires the most extra hardware. Although MetroHA is best suited for use cases up to 256 terabytes, the average deployment is between 32 and 128 terabytes per the MetroHA cluster. Auto-Sync is part of Nexenta's HA Cluster licensed feature. It is a schedulable, fault-managed, fully configurable data replication service that can be used in a variety of backup, archiving, and disaster recovery scenarios, including offsite data compliance, online data migration, websites, and back office systems.

The beauty of the Nexenta solution is that it's all software-defined. Because it's a software layer that is leveraging commodity hardware, customers get better utilization, better orchestration, and CapEx and OpEx savings.

The Bigger Truth

More data has been created in the last three years than in all of the past 40,000 years. Now is the time for a new storage strategy to deal with the ongoing avalanche of big data and unstructured data, as well as the demands for the increasingly important five 9s availability of this data. The old hardware strategy was strictly one dimensional—when in doubt, throw big iron at storage problems. The evolution to software-defined storage with commodity hardware components shifts the storage management equation; it allows enterprises to focus on managing and serving data rather than hardware. It's true that not every bit of data is suitable for the SDS, commodity-based hardware model. Sensitive tier-1 financial data, for instance, is likely to still be housed on the big iron, scale-up hardware that it has always resided on. But much of today's unstructured data is well suited for the new, commodity-based SDS systems.

Big iron architectures—and the big price tags they come with—were designed for five 9s availability. However, the costs of protecting against local failures—hardware, software, site, or human—and regional or global disasters were enormous and prohibitive. Thanks to software-defined storage, today's IT organizations can achieve the same levels of availability using commodity hardware. It brings five 9s availability to companies that could never before imagine attaining those levels because, before now, they couldn't afford it.

Nexenta is a prominent player in the software-defined storage market that provides affordable five 9s availability on commodity hardware. Its flagship product, NexentaStor, offers a scale-out, node-based architecture managed by software-defined storage at the top layer. Nexenta's customers are able to scale up or scale down storage and IOPS, depending on business flows, using the best hardware for the available budget.

Big data is a big opportunity, but it also leads to big questions. Software-defined storage can definitely answer some of them—especially for managing data availability.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com